

SVC-L2.5 V PNs: COMBINING LAYER-3 V PNs TECHNOLOGY WITH SWITCHED MPLS/IP L2VPNs FOR ETHERNET, ATM AND FRAME RELAY CIRCUITS

RELATED U.S. APPLICATION DATA

[0001] Provisional application No. 60/409,325 filed on September 9, 2002.

FIELD OF THE INVENTION

[0002] The present invention relates to a combination of switched Layer-2 Virtual Private Networks (VPNs) with a layer-3 VPN and is particularly concerned with flexible, on-demand switched MPLS/IP Layer-2 VPNs for Ethernet, ATM and Frame Relay SVCs while distributing customer routes through Layer-3 VPN mechanisms.

BACKGROUND OF THE INVENTION

[0003] A Virtual Private Network (VPN) may be thought of as a private network constructed within a shared network infrastructure. In common terminology, these private networks are used by clients while the network infrastructure is supplied by providers.

[0004] Existing varieties of Layer-3 VPNs have limitations affecting ease of implementation and use generating:

- [0005]** - customers who are not comfortable with Layer-3 VPN IP datapath due to security concerns;
- [0006]** - customers who want to have flexibility to use Layer-2 circuits for some applications integrated into a Layer-3 VPN;
- [0007]** - customers who want to improve their Layer-3 management but want control on the datapath;
- [0008]** - customers who want to use Layer-2 QoS capabilities with IP-VPN service; and

[0009] - customers who plan to upgrade to Layer-3 VPN in the future but are not ready to give up their existing Layer-2 networks, for example, Frame Relay networks currently generating revenue.

[0010] In view of the foregoing, it would be desirable to provide a technique for providing switched layer-2 VPNs combined with a subset of layer-3 VPN technology which overcomes the above-described inadequacies and shortcomings.

SUMMARY OF THE INVENTION

[0011] An object of the present invention is to provide an improved switched virtual circuit Layer-2 virtual private network arrangement combining Layer-3 VPNs technology with switched MPLS/IP L2VPNs for Ethernet, ATM and Frame Relay Circuits.

[0012] According to an aspect of the present invention, [there] is provided a network for providing switched virtual circuit Layer-2 VPNs, wherein the network includes a set of elements interconnected by services; at least one first subset of said elements defining a private network; and at least one second subset of elements different from said first subset defining a provider network wherein at least two subgroups of said first subset of elements may be connected via said provider network. There are a plurality of customer ports maintained on the elements of the first subset of elements and a plurality of provider ports maintained on the second set of elements, each of the plurality of provider ports connected by services to a customer port. At each element of the provider network having a provider port is a port information table containing mapping information relating addresses of customer ports to addresses of provider ports for the first subset of elements. The network also includes a provisioning mechanism used to define element membership in said first subset of elements, a signalling mechanism used to create Layer-2 connectivity between elements within said first subset of elements at the Layer-2 level across said second subset of elements, and a reachability distribution mechanism.

[0013] Advantages of the present invention include real-time establishment of customer Layer-2 virtual circuits (VCs), and the ability to perform dynamic client reconfiguration via dynamic routing. Support for traffic engineering within the L2.5VPN service can be rendered without impacting traffic engineering on the provider network. There is support for an arbitrary mesh topology. In terms of mobility, L2.5VPN allows the ability to move one port of an L2.5VPN from one provider edge device (PE) to another and one provider to another without impacting the L2.5VPN and client network addressing. A further advantage is that L2.5 dynamic bandwidth management supports interworking to legacy Layer-2 VPNs.

[0014] Conveniently the invention further provides for the reachability distribution mechanism to use a Layer-3 VPN service. This Layer-3 VPN service could be one of piggybacking VPN routes onto the backbone Border Gateway Protocol, or alternatively that of using a virtual router redistribution scheme

[0015] Conveniently the invention further provides for an auto-discovery mechanism for distributing said mapping information to layer-2 port information tables of the provider network. This auto-discovery mechanism for distributing said mapping information uses Border Gateway Protocol in some instances.

[0016] In accordance with another aspect of the present invention, there is provided a method of organizing a network having a set of elements interconnected by services, wherein at least one first subset of the elements defines a private network and at least one second subset of elements different from the first subset defines a provider network and wherein at least two subgroups of the first subset of elements may be connected via the provider network, wherein the method includes the steps of defining element membership in said first subset of elements via a provisioning mechanism; establishing a plurality of customer ports within said elements of the first subset of elements; and establishing a plurality of provider ports within the second set of elements. Each of the plurality of provider ports are connected by services to a customer port. Thereafter, the step of establishing a port information table at each element of said provider network having a provider port, the port information table

containing mapping information relating addresses of customer ports to addresses of provider ports. The method further includes the steps of determining reachability across said second subset of elements; and creating Layer-2 connectivity within the first subset of elements at the Layer-2 level across the second subset of elements via a signalling mechanism.

[0017] The present invention further includes a method of organizing a network having a set of elements interconnected by services, wherein at least one first subset of the elements defines a private network and at least one second subset of elements different from the first subset defines a provider network and wherein at least two subgroups of the first subset of elements may be connected via the provider network. The method includes the steps of defining a L2VPN topology; establishing a plurality of customer ports within said elements of said first subset of elements; and establishing a plurality of provider ports within said second set of elements, each of said plurality of provider ports connected by data and signalling services to a customer port. Thereafter, creating a Layer-2 Port Information Table for each provider port; establishing the identity of customer ports attached to each provider port, and populating the Layer-2 Port Information Table at that provider port with mapping information relating addresses of customer ports to addresses of provider ports. Further steps include distributing said mapping information to Layer-2 Port Information Tables of the provider network via an auto-discovery mechanism; determining reachability across the second subset of elements via a Layer-3 VPN service; and creating Layer-2 connectivity within the first subset of elements at the Layer-2 level across the second subset of elements via a signalling mechanism upon request from an element within the first subset of elements.

[0018] The present invention will now be described in more detail with reference to exemplary embodiments thereof as shown in the appended drawings. While the present invention is described below with reference to the preferred embodiments, it should be understood that the present invention is not limited thereto. Those of ordinary skill in the art having access to the teachings herein will recognize additional

implementations, modifications, and embodiments which are within the scope of the present invention as disclosed and claimed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0019] The invention will be further understood from the following detailed description of embodiments of the invention and accompanying drawings in which:
- [0020] **FIG. 1** is a diagram of a generic network having a shared network infrastructure and Virtual Private Networks associated thereto;
- [0021] **FIG. 2** is a diagram of a network reference model including a plurality of customer edge devices, provider edge devices, and provider devices within the network;
- [0022] **FIG. 3** is a diagram of the relation between Layer-2 datapath and network services provisioned by the service provider according to an embodiment of the invention;
- [0023] **FIG. 4** is a diagram of a L2.5VPN network according to an embodiment of the invention;
- [0024] **FIG. 5** is a block diagram of SVC-L2.5VPN mechanisms according to an embodiment of the invention;
- [0025] **FIG. 6** is a diagram of a L2.5VPN network depicting one version of Layer-3 reachability distribution according to an embodiment of the invention; and
- [0026] **FIG. 7** is a diagram of a L2.5VPN network depicting another version of Layer-3 reachability distribution according to an alternative embodiment of the invention.

DETAILED DESCRIPTION

[0027] **Glossary of Acronyms Used**

P – Provider Device
PE – Provider Edge Device
CE – Customer Edge Device
SVC – Switched Virtual Circuit
CPI – Customer Port Identifier (Layer-2)
PPI – Provider Port Identifier (Layer-2)
PIT – Port Information Table
L2PIT – Layer-2 Port Information Table
BGP – Border Gateway Protocol
BGP-AD – BGP Auto-Discovery
MPLS – Multi-Protocol Label Switching
DLCI – Data Link Connection Identifier
LMP – Link Management Protocol
ISP – Internet Service Provider
SVC-TE – SVC-L2VPN with Traffic Engineering Capabilities

[0028] Referring to FIG. 1, there may be seen a generic network having a shared network infrastructure 100 with connected virtual private network sites 101. The VPN sites 101 make use of the network infrastructure 100 to interconnect physically remote sub-networks of particular VPNs.

[0029] Referring to FIG. 2, there may be seen a network reference model showing a more detailed depiction of a network having a plurality of customer edge router/switches (CEs) 201, 202, 203, 204, 205, 206, 207, 208 and 209. The provider network has provider edge router/Layer-2 switches (PEs) 210, 212, and 214 as well as provider devices (P) 215, 216, 217, and 218 interior to the provider network.

[0030] Further in FIG. 2 may be seen the typical case where VPN A has a portion connected to CEs 201 and 202, and another portion connected to CE 206. Communication services between these remote portions of VPN A will be provided by the provider network. The same general situation obtains for VPN B, VPN C, and VPN D.

[0031] In operation, the Switched Virtual Connection Layer-2.5 VPN (SVC-L2.5VPN) is a provider-based Layer-2 and Layer-3 VPN service that allows clients to request on-demand Layer-2 circuits while distributing customer routes through Layer-3 mechanisms.

[0032] The SVC-L2.5VPN uses the mechanisms for SVC-L2VPN described in US Patent Application _____, hereby incorporated by reference, which are characterized by:

[0033] – a given topology;

[0034] – using IP/MPLS based signalling between CE-PE (or any other layer-2 signalling protocols);

[0035] – the possible employment of Link management protocol (LMP) for Layer-2 link-port consistency;

[0036] – use of private addresses which have the potential to be overlapping with other addresses in other VPNs; and

[0037] – the capacity to be built using single-sided signalling and auto-discovery mechanisms as, for example, being standardized in IETF.

[0038] Layer-2.5 VPN service combines both advantages of Layer-3 VPNs as described in RFC2547 and “switched” Layer-2 VPNs in that:

[0039] – it allows the CE to peer with the PE at Layer-3 only i.e. there is no need to peer with all remote CEs;

[0040] – it allows the CE to use a Layer-2 VPN as the transport mechanism; and

[0041] – it also allows for re-using the advantages of new GMPLS-enabled VPN technology, namely to separate datapath from control, and to perform single-ended provisioning.

[0042] A formulaic description would be as follows:

SVC MPLS/IP L2.5VPN = SVC + (G)MPLS + IP + VPN Constructs

[0043] where:

SVC implements the private switched model;

(G)MPLS provides signalling for Layer-2 connections;

IP is the IP control channel and IP VPN route distribution; and

VPN Constructs are services such as VPN membership, overlapping addresses, VPN auto-discovery, etc.

[0044] The key objectives of Layer-2 use in L2.5VPNs includes:

[0045] – constrained or restricted connectivity as defined by customer, and as maintained and enforced by the service provider;

[0046] – an on-demand Layer-2 circuit request initiated by the L2.5VPN customer requiring no coordination with the service provider;

[0047] – the client devices operate within the L2.5VPN space independently from the service provider network operations Subject to the defined constrained or restricted connectivity;

[0048] – there exists privacy/independence with respect to addressing and routing both among L2.5VPN customers, as well as between an L2.5VPN customer and a service provider;

[0049] – there is support for single-ended provisioning; and

[0050] – there is support for a multiservice Layer-2 switched model including such services as ATM, Frame Relay, Ethernet, Ethernet VLAN (PPP, HDLC, etc.).

[0051] The key objectives of Layer-3 use in L2.5VPNs includes:

[0052] – Layer-3 VPN constructs, specifically distributing reachability using VPN distribution of VPN routes through the backbone BGP (as per RFC2547), or virtual router (VR) distribution of VPN routes; and

[0053] – optional IP services for IP traffic (if a L2.5VPN is provisioned to also provide a layer-3 VPN).

[0054] A number of benefits for both client and provider are associated with SVC-L2.5VPNs as compared to legacy Layer-2 VPNs and SVC-L2VPNs.

[0055] Advantages to the VPN Customer on the client side are multiple and include:

[0056] - no peering with private sites (resolving n-square routing peering issue);

[0057] - peering only with attached PE;

[0058] - being able to use Layer-2 Circuits even when an L3VPN is offered, thereby taking advantage of legacy and new Layer-2 VPNs;

[0059] - compatible with access clients that are 'MPLS/IP' signalling based

[0060] - supports overlapping/private address space;

[0061] - supports Layer-3 addresses within the L2VPN (and does not require transport Layer-2 addresses);

[0062] - higher mobility in that a customer can move its L2VPN from one port to another without changing the addressing of the L2VPN (in fact, without changing the L2VPN addressing, QoS, etc.) thus offering a greater flexibility for network operations;

[0063] - that the L2VPN addresses can be used for customer Layer-3 network;

[0064] - offering a range of security capabilities including Layer-2 security;

[0065] - offering a range of QoS capabilities that includes Layer-2 VPNs QoS (including the legacy L2VPNs);

[0066] - allowing the SVC-L2VPN circuit to be used as either a legacy Layer-2 circuit or as an MPLS LSP within the client network as needed;

[0067] - not requiring the client to implement full MPLS but just signalling protocol at the edges; and

[0068] - allowing the option of the client using the SVC-TE services to better optimize his network and perform traffic engineering operations.

[0069] Advantages to the Service Provider on the provider side include:

[0070] - opportunity for new revenue opportunities to the ISPs;

[0071] - support for Dynamic Membership distribution to ease circuit configuration and distribution;

- [0072] - capable of interworking with existing legacy Layer-2 VPNs;
- [0073] - provides opportunity to maximize yield from network investment on legacy Layer-2 and IP/MPLS based infrastructure;
- [0074] - leverages existing provider skill level in Layer-2 VPNs;
- [0075] - avoids requirement for tunnelling (including MPLS) between PE-PE (only when MPLS is used in the core);
- [0076] - support for reusing (G)MPLS for link, port constructs ;
- [0077] - support for single-sided signalling;
- [0078] - allows Provider network operations to be completely decoupled from the customer L2VPNs unlike the case for legacy switched L2VPNs; and
- [0079] - provides better scaling than Layer-3 VPNs or Layer-3 VPNs with extended two-phase discovery mechanisms.
- [0080] Dependent upon the implementation and service offering, an L2.5VPN service can offer:
 - [0081] – a L2.5VPN service with options to offer an Layer-3 VPN service (in addition to L2.5) on the same port if needed.
 - [0082] – to use an Layer-2 VPN service (in addition to L2.5) consisting of:
 - [0083] - traditional legacy L2VPN;
 - [0084] - new MPLS/IP L2VPN (PVC models);
 - [0085] - new MPLS/IP Switched L2VPN;
 - [0086] - new SVC-TE (L2VPN with traffic engineering capabilities);
 - [0087] – to exclusively operate a L2.5VPN service but with Layer-3 reachability distribution, and Layer-2 datapath; and
 - [0088] – a L2.5VPN service with traffic engineering (TE) capabilities.
- [0089] The Layer-3 advantages of a L2.5VPN service include:
 - [0090] - IP access between PE and CE at control plane;
 - [0091] - the CEs do not establish routing peering between themselves;
 - [0092] - the PE devices manage customer routes for distribution only;
 - [0093] - the service provider (SP) network provides automatic inter-site connectivity among customer CE devices;

[0094] - the SP guarantees security and isolation of the VPNs between themselves and between the service provider's network(s) using a variety of options including that of legacy L2VPNs; and

[0095] the SP may offer per VPN basis extranet and internet access with an L2.5VPN.

[0096] The SVC-L2.5VPN protocol requirements are as follows:

[0097] - at the CE:

[0098] - require support for MPLS signalling, for example RSVP-TE with SVC-L2VPN extensions but not necessarily MPLS forwarding;

[0099] - require an IP-based control channel, for example, IP tunnelling; and

[0100] - require support for routing-protocol or static routes between CE-PE within the VPN context.

[0101] - at the PE:

[0102] - require an IP based control channel;

[0103] - require MPLS signalling; and

[0104] - optionally an auto-discovery mechanism; and

[0105] - require a mechanism for distributing reachability

[0106] The SVC-L2.5VPN Architecture Components may be summarized as follows:

[0107] - Access is Layer-2 or Layer-3 VPNs;

[0108] - Require an IP-based control channel for learning customer routes and signalling adjacency

[0109] - Layer-2.5 VPN reachability distribution using a Layer-3 VPN service for distributing reachability, such as piggybacking VPN routes onto the backbone BGP as described in RFC2547, or by using a Virtual Router (VR) distribution scheme;

[0110] - a generalized L2.5 discovery mechanism using Layer-3 discovery for the Layer-3 routes, and Layer-2 discovery for the Layer-2 port information;

[0111] - Membership is defined in the same way as existing Layer-2 VPNs and Layer-3 VPNs;

- [0112] - switched on-demand SVC-L2VPNs;
- [0113] - forward adjacency with L2.5VPNs;
- [0114] – Ports and links are logical constructs that uses (G)MPLS functions; and
- [0115] – Signalling is MPLS based (packet side only) between CE-PE.
- [0116] The SVC-L2.5VPN Building Blocks may be summarized as follows:
- [0117] – Customer and Provider Ports;
- [0118] – A Layer-2 Port Information Table (L2PIT) which maintains mapping between customer ports and provider ports (at the edges of the service provider network) provides local CEs with the information about other ports in the SVC-L2.5VPN, and is defined on a per SVC-L2.5VPN basis or for all the SVC-L23VPNs connected to PE;
- [0119] – a Layer-2 BGP based auto-discovery mechanism (BGP-AD) used to determine and distribute information related to customer and provider ports to the PEs, and to populate the L2PIT with this information;
- [0120] – a (G)MPLS-signalling mechanism to create connectivity within the set of client devices that are part of the same VPN at the Layer-2 level; and
- [0121] – a reachability distribution mechanism which may be VPN distribution of VPN routes through the backbone BGP, or virtual router (VR) distribution of VPN routes.
- [0122] Customer site reachability may be determined either by:
- [0123] – use of static routes; or
- [0124] – use of standard routing protocols such as RIP, OSPF, or IBGP.
- [0125] Referring to **FIG. 3**, there is depicted a representation of the relationship between a Customer Edge device 301 (CE), typically a router; the Layer-2 datapath 303; the private routes 305 defined separately from the datapath; the service provider network 307; and the provider provisioned Layer-2.5 VPN architecture layers 309.
- [0126] Referring to **FIG. 4** a L2.5VPN may be seen having a provider network with backbone 401, and provider edge device 403 (PE). A customer edge device 405 (CE) connects via Layer-2/Layer-3 access 407 to the provider edge device 403.

Between edge devices **403** and **405** run services **409** consisting of OSPF/RIP/BGP, and/or MPLS signalling for L2VPN. The Layer-2 Virtual circuit **411** connects remote sections of VPN A through the provider network. The reachability distribution is Layer-3 VPNs, and the datapath is Layer-2 VPNs. With L2.5VPNs with traffic engineering (L2.5VPN-TE), the CE **405** will form a forwarding adjacency out of that Switched Virtual Circuit (SVC) by advertising the SVC as a TE link into the same instance of ISIS/OSPF. The SVC-L2VPN circuit can appear as an MPLS LSP to the CE **405** if the CE **405** is running MPLS.

[0127] The use of L2.5VPN provides for simplified provisioning in that:

[0128] – Addition of a new port to a given SVC-L2.5VPN involves configuration and/or provisioning changes only on the PE that has this port;

[0129] – BGP is used to distribute this information to other PEs that have ports of that SVC-L2VPN;

[0130] – BGP is used to distribute this information to other CEs that have ports of that SVC-L2VPN;

[0131] – the customer could establish or terminate a Layer-2 connection between a pair of ports in its SVC-L2VPN without involving configuration or provisioning changes in any of the service provider equipment by using (G)MPLS signalling; and

[0132] – the customer establishes a Layer-3 peering only with the attached PE.

[0133] The SVC-L2.5VPN mechanisms are illustrated in **FIG. 5** where the L2.5VPN **501** has two distinct operations: the Switched L2VPN operations **502**, and the Layer-3 VPN operations **503**. Subsumed under the Switched L2VPN operations **502** are the learning customer port information **504** and Port Information Table build out and Port Information distribution **506**. Subsumed under the Layer-3 VPN operations **503** are the learning customer reachability information function **505** and the Layer-3 distribution phase function **507**. Switched L2VPN Operations **502** can offer offline Traffic Engineering as an option. Switched L2VPN Operations **502** can also be accomplished by using GMPS based optical VPNs. Layer-3 VPN operations **503** are only for online Traffic Engineering, which is accomplished by Layer-3 VPNs mechanisms.

[0134] Referring to **FIG. 6**, there is a depiction of a L2.5VPN with a reachability distribution scheme consisting of BGP updates through the backbone BGP. The backbone **601** connects a pair of Provider Edge devices **603**, typically routers. A Virtual Router **605** connects to VPN A, while VPN Reachability Information **607** traverses backbone **601** via BGP updates **609**. This is an illustration of the first of the Layer-3 reachability distribution methods.

[0135] By way of contrast, **FIG. 7** illustrates a L2.5VPN with a Virtual Router reachability distribution. VPN A **701** connects to Provider Edge device **703**, typically a router, which connects across the provider network to a second Provider Edge device **705**. Virtual Routers **704** and **706** provide the virtual routing mechanism within the Provider Edge devices. Item **709** represents a routing instance, and items **711** show routing updates providing per VPN reachability information along tunnels **707** which run along the backbones **713** of the provider network. This is an illustration of the second of the alternate Layer-3 reachability distribution methods.

[0136] While the invention has been described in conjunction with specific embodiments thereof, it is evident that many alternatives, modifications, and variations will be apparent to those skilled in the art in light of the foregoing description. Accordingly, it is intended to embrace all modifications, variations and adaptations such as may be made to the particular embodiments of the invention described above without departing from the scope of the invention, which is defined in the claims.